arXiv:0709.4529v1 [math-ph] 28 Sep 2007

ON THE NEIGHBOR SPACING OF EIGENVALUES OF UNITARY MATRICES

DAVID W. FARMER

Abstract. We describe a subtle error which can appear in numerical calculations involving the spacing statistics of eigenvalues of random unitary matrices.

1. Introduction

The Random Matrix Theory (RMT) of the classical compact groups has become a central tool in analytic number theory. By modeling the Riemann zeta-function and other L-functions with the characteristic polynomial of a random unitary matrix, new conjectures have been made about L-functions, which have also been useful for proving new results. One of the reasons for this success is that many quantities which are mysterious in number theory can be computed exactly in RMT. Examples include the statistical behavior of the zeros [6, 7, 3, 8], value distribution [4, 5], and moments [1].

There are some questions about L-functions for which the RMT analogue has not yet been answered. But in many of those cases one can exploit the fact that it is easy to generate Haar-distributed random matrices for the unitary, symplectic, and orthogonal groups. See Mezzadri [9] for a complete discussion. Thus, one can usually obtain numerical evidence to support a conjecture or to check a calculation.

In many cases, data generated from relatively small matrices is almost indistinguishable from the limiting case of large matrices. For example, the normalized nearest neighbor spacing of eigenvalues of 15×15 random unitary matrices differs from the limiting case of large unitary matrices by so little that the difference is not readily visible in a histogram.

Despite the simplicity of generating random matrices, it turns out there is a subtle bias in the way many people generate and perform experiments on random matrices. I have committed this error (not realizing it was an error) many times, and it never made a difference. But then I tried to check the next-to-leading order term of a fairly involved RMT calculation [2]. The numerics refused to confirm the calculation, and it turned out that the numerics were wrong. Fortunately, the error is easily avoided once you know about it.

The next section describes a simple example of the situation in which this error occurs, and gives some data to illustrate the problem. In section 3 the underlying cause of the problem is explained.

2. The neighbor spacing of eigenvalues

The sample problem concerns the neighbor spacing of the eigenvalues of random unitary matrices. These spacings are completely understood: we choose this example to illustrate what can go wrong in the numerical experiments.

Research supported by the American Institute of Mathematics and the NSF focused research group grant DMS0244660.

Suppose $U \in U(M)$ is an $M \times M$ unitary matrix. The characteristic polynomial of U is

$$\Lambda_U(x) = \prod_{m=1}^{M} (x - e^{i\theta_m}) \tag{2.1}$$

where $\{e^{i\theta_1}, \ldots, e^{i\theta_M}\}$ are the eigenvalues of U and $-\pi \le \theta_1 \le \cdots \le \theta_M < \pi$ are the *eigenangles* of U. Let $\delta_m = \frac{M}{2\pi}(\theta_{m+1} - \theta_m)$ denote the normalized neighbor spacing of the eigenangles, where we set $\theta_{M+1} = \theta_1 + 2\pi$ so that δ_M is the "wrap around" neighbor spacing. Collectively the M numbers $\delta_1, \ldots, \delta_M$ are 1 on average, because $\delta_1 + \cdots + \delta_M = M$ by construction. but as commonly implemented on a computer, the individual δ_j are *not* one on average. That is the point of this paper.

2.1. **Generating random matrices.** The computer code and data in this paper are from Mathematica, but the code should be understandable without in-depth knowledge of Mathematica, and the apparent anomalies in the data have nothing to do with Mathematica.

The following code defines a function `randunitary[M]` which generates a Haar-random matrix in $U(m)$. The first line loads the Mathematica package required to generate normally distributed random numbers. The second line sets up the generation of the normally distributed random numbers, and the third command implements the algorithm described by Mezzadri [9].

```
<<Statistics‘NormalDistribution‘

norm = NormalDistribution[0, 1]

randunitary[M_]:=Block[{gmatrix,q,r,d,dd},
    gmatrix=Table[Random[norm]+Random[norm]*I,{i,1,M},{j,1,M}]];
    {q,r} = QRDecomposition[gmatrix];
    d=DiagonalMatrix[Table[dd=r[[j,j]];dd/Abs[dd],{j,1,M}]];
    q.d]
```

Note that this code already fixes a non-obvious error that is present in the most straightforward way to generate a random unitary matrix. Namely, the matrix `q` in the above code should be a Haar-random matrix in $U(M)$, but it isn't. The problem is that Mathematica's `QRDecomposition` function introduces a bias which causes the matrix `q` to have fewer eigenvalues very close to 1. Multiplying by the diagonal matrix `d` fixes that problem. This issue arises in the `QRDecomposition` routine of every computer algebra package. See Mezzadri [9] for details.

Note also that this code is for Mathematica 5. Some modifications may be needed for Mathematica 6 because the functions for generating random numbers have been changed.

The following code extracts the eigenangles and defines the normalized neighbor spacings. In this example we use 14×14 matrices.

```
M=14;
mymatrix=randunitary[M];
eigangles= (M/(2 Pi)) Sort[Arg[Eigenvalues[mymatrix]]];
delta[j_]:=eigenangles[[j+1]] - eigenangles[[j]];
delta[M]:=eigenangles[[1]] - eigenangles[[M]] + M;
```

2.2. **Problems with the eigenangle spacing.** Depending on how you use it, the above code has an error. Here is how the error manifests itself. For fixed j, the average (expected) value of `delta[j]` is *not equal to 1, even in the large M limit*. While it is true that, by construction, the average of `delta[1],...,delta[M]` is identically equal to 1, each individual `delta[j]` does not average to 1. Table 2.1 shows the average of `delta[j]` for various j and M, averaging over 100,000 matrices.

$M =$	14	22	32
δ_1	0.94345	0.94597	0.94506
δ_3	0.99367	0.99549	0.99387
δ_7	0.99912	0.99836	1.00045
δ_{11}	0.99352	0.99926	0.99745
δ_{rand}	1.00260	0.99948	1.00147

TABLE 2.1. Averages of δ_j for random matrices in $U(M)$ for various j and M. Each column is the average for 100,000 matrices. The bottom row averages δ_j for a randomly chosen j for each matrix.

As the table shows, at least some of the selected δ_j differ significantly from 1. Curiously, most are less than 1 on average. The last row, δ_{rand}, is generated by randomly using one of $\delta_1, \ldots, \delta_M$ for each matrix. and then averaging those random choices. We know that δ_{rand} must average to 1, and the closeness of δ_{rand} to 1 should suggest that (at least some of) the other values differ significantly from 1, as opposed to arising from fluctuations due to a small sample size.

2.3. **More curiosities in the neighbor spacings.** The example in the previous section is unrealistic, because nobody would use only one neighbor spacing for each matrix. Even if that gave the right answer, it would be much more efficient to use several, or all, of the available neighbor spacings. However, it is quite common for people to just choose the "lazy person's neighbor spacings" $\delta_1, \ldots, \delta_{M-1}$. That is, throw away the "wrap around" spacing, which avoids the need to program that as a special case. One could argue (incorrectly!) that since Haar measure is rotationally invariant, all the neighbor spacings are the same, and so no bias is introduced by omitting one value. That argument is wrong, as we shall see. Here is the average of $\delta_1, \ldots, \delta_{M-1}$ for the same data set used to generate Table 2.1.

$M =$	14	22	32
average of $\delta_1, \ldots, \delta_{M-1}$	0.9862	0.99157	0.99419
average of δ_M	1.1796	1.1768	1.17988

TABLE 2.2. Average of $\delta_1, \ldots, \delta_{M-1}$, and the average of δ_M, for the same data set of 100,000 matrices used to generate Table 2.1.

As Table 2.2 shows, the average of the "lazy person's neighbor spacings" is less than 1. If that is true, then the average of the "wrap around" spacing must be greater than 1, which Table 2.2 also confirms.

The fact that δ_M is larger than average is not due to an error in the code. It is due to a bias caused by the way we choose the eigenangles, which we describe in the next section.

3. How not to choose a random gap

Suppose you want to choose a random gap between the eigenangles of a random matrix, so that each gap is equally likely. The following procedure is incorrect: choose a random point on the unit circle, and then select the gap which includes that selected point. That method introduces a bias, because large gaps are more likely to be selected than small gaps. For example, if all the eigenvalues were in one-half of the unit circle, the biggest gap would be selected at least half the time. That is, the expected size of a gap surrounding any given point will be larger than average.

That bias is implicit in the way the eigenangles have been chosen in the Mathematica code above. Because we restrict the eigenangles to the interval $[-\pi, \pi)$, the "wrap around" gap is chosen to include the point -1. Thus, the expected size of δ_M should be larger than average, as the data shows. A calculation finds that the expected size of a gap selected in this way is the mean of the *square* of the nearest neighbor spacing. For the large M limit of eigenangles from $U(M)$, the square of the normalized nearest neighbor spacing is approximately 1.180, which agrees well with the data in Table 2.2. And since g_M is larger than average, the rigidity in the eigenangle spacing forces g_1 and g_{M-1} to be smaller than average. The effect fades away so that $g_{M/2}$ is only slightly smaller than average, That is, for fixed j, the expected value of δ_j is less than 1, although the amount less than 1 is a rapidly decreasing function of j. The average of $\delta_1, \ldots, \delta_{M-1}$ is approximately $1 - 0.18/M$, but even that can be a significant difference if one is exploring the dependence on the size of the matrix.

Thus, even if the distribution of points on the unit circle is rotationally invariant, there is a bias introduced by the simple procedure of taking the `Arg[]` of the points and dealing instead with numbers in the interval $[-\pi, \pi)$. In particular, the shortcut of leaving out the "wrap around" gap can have unintended consequences.

References

[1] J.B. Conrey, D.W. Farmer, J.P. Keating, M.O. Rubinstein, and N.C. Snaith, *Integral moments of L-functions*, Proceedings of the London Mathematical Society, **91**, 33–104.

[2] E. Dueñez, D. W. Farmer, S. Froehlich, C. P. Hughes, F. Mezzadri, and T. Phan, *Roots of the derivative of the Riemann zeta function and of characteristic polynomials*. preprint.

[3] N. M. Katz and P. Sarnak, Random matrices, Frobenius eigenvalues, and monodromy, AMS Colloquium Publications, 45, AMS, Providence, RI (1999).

[4] J. P. Keating and N. C. Snaith, *Random matrix theory and* $\zeta(\frac{1}{2} + it)$, Comm. Math. Phys. **214** (2000) pp. 57–89.

[5] J. P. Keating and N. C. Snaith, *Random matrix theory and L-functions at* $s = \frac{1}{2}$, Comm. Math. Phys. **214** (2000) pp. 91–110.

[6] H.L. Montgomery, *The pair correlation of zeros of the Riemann zeta-function*, Proc. Symp. Pure Math. **24** (1973) pp. 181-93.

[7] A. Odlyzko, *The* 10^{20}*th zero of the Riemann zeta-function and 70 million of its neighbors.*, Preprint (1989).

[8] M. O. Rubinstein, *Evidence for a spectral interpretation of the zeros of L-functions*, (1998) thesis, Princeton University.

[9] F. Mezzadri *How to generate random matrices from the classical compact groups*, To appear in Notices of the AMS, math-ph/0609050.

American Institute of Mathematics
farmer@aimath.org